

# Parallel Adaptive High-Order CFD Simulations Characterizing Cavity Acoustics for the Complete SOFIA Aircraft

Michael F. Barad  
NASA Ames Research Center  
Moffett Field, CA 94035  
michael.f.barad@nasa.gov

Christoph Brehm  
Science and Technology Corp.  
Moffett Field, CA 94035  
christoph.brehm@nasa.gov

Cetin C. Kiris  
NASA Ames Research Center  
Moffett Field, CA 94035  
cetin.c.kiris@nasa.gov

Rupak Biswas  
NASA Ames Research Center  
Moffett Field, CA 94035  
rupak.biswas@nasa.gov

## ABSTRACT

This paper presents one-of-a-kind MPI-parallel computational fluid dynamics simulations for the Stratospheric Observatory for Infrared Astronomy (SOFIA). SOFIA is an airborne, 2.5-meter infrared telescope mounted in an open cavity in the aft of a Boeing 747SP. These simulations focus on how the unsteady flow field inside and over the cavity interferes with the optical path and mounting of the telescope. A temporally fourth-order Runge-Kutta, and spatially fifth-order WENO-5Z scheme was used to perform implicit large eddy simulations. An immersed boundary method provides automated gridding for complex geometries and natural coupling to a block-structured Cartesian adaptive mesh refinement framework. Strong scaling studies using NASA's Pleiades supercomputer with up to 32,000 cores and 4 billion cells shows excellent scaling. Dynamic load balancing based on execution time on individual AMR blocks addresses irregularities caused by the highly complex geometry. Limits to scaling beyond 32K cores are identified, and targeted code optimizations are discussed.

## Categories and Subject Descriptors

J.2 [Computer Applications]: Physical Sciences and Engineering—*Aerospace*; I.6 [Computing Methodologies]: Simulation and Modeling

## 1. INTRODUCTION

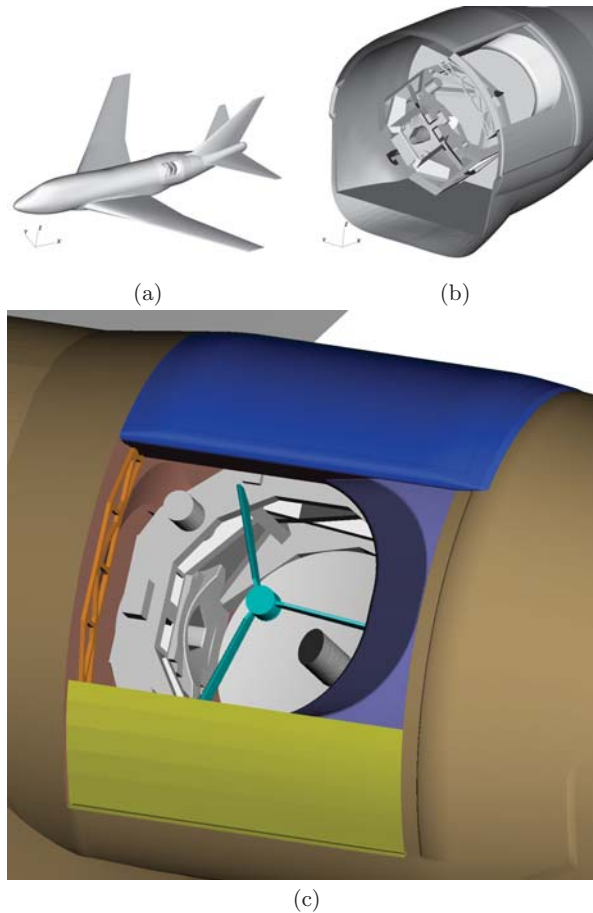
This paper presents MPI-parallel computational fluid dynamics (CFD) simulations for the Stratospheric Observatory for Infrared Astronomy (SOFIA). SOFIA is an airborne, 2.5-meter infrared telescope mounted in a large open cavity in the aft fuselage of a Boeing 747SP, as shown in Figure 1. The aircraft typically flies high in the atmosphere, between 12-



Figure 1: SOFIA aircraft and aft open cavity where a 2.5-meter infrared telescope is mounted.

14 kilometers (39,000 to 45,000 feet), reducing atmospheric distortion and enabling high quality images. SOFIA is a joint program between NASA and the German Aerospace Center (DLR). In this work, we focus on 1) noise generation and flow physics in the open cavity where the telescope is mounted, and 2) assessment of solver parallel performance for up to 32K cores on this multi-scale, complex geometry, demanding application.

Higher-order accurate implicit large eddy simulations (ILES) using NASA's Launch Ascent and Vehicle Aerodynamics (LAVA) CFD solver [14] were performed. Direct numerical simulation, and for realistic Reynolds numbers, more practically relevant large eddy simulations are an essential tool to understand the complex flow physics of unsteady flow fields. These simulations become computationally very expensive due to the requirement of accurately capturing a wide range of physically and temporally relevant scales. Numerical methods with high spectral resolution are required to efficiently simulate these types of problems. The computation of transonic flows is further complicated by the occurrence of shocks which generate discontinuities in the flow field. In order to handle these discontinuities, we apply state-of-the-art higher-order shock capturing schemes. The different schemes available within LAVA were analyzed with respect to accuracy and robustness in Brehm *et al.* [5]. LAVA's immersed boundary method was used to circumvent the very time-consuming volume mesh generation process for the complex SOFIA geometry. The current ghost cell based immersed boundary method is a natural fit with the block-



**Figure 2: Geometry used for SOFIA 747SP, including detailed open cavity with telescope hardware. (a) modeled aircraft, (b) truncated view of cavity and telescope, and (c) view of open cavity from outside.**

structured Cartesian adaptive mesh refinement framework.

In previous work, two critical issues with respect to SOFIA telescope imaging quality were identified [11]. Firstly, the local refraction indices within the line of sight are strongly affected by the spatial density distribution in the shear layer. High fidelity CFD simulations provide accurate density distributions which can be used to calculate the effective optical path length and improve the image quality. Secondly, the unsteady pressure field inside the cavity can induce telescope vibrations. Targeted structural damping can be applied to the most dominant frequencies identified by CFD.

The current simulations focused on how the unsteady flow field inside and over the open cavity interferes with the optical path and mounting of the telescope. Results include the identification of two important mechanisms potentially affecting image clarity: (1) density variations in the shear layer leading to image distortion, and (2) vibrations of the telescope structure perturbing the optical path, hence, reducing image quality. The “secondary” breakdown process

of the shear layer involves the interaction and generation of vortical flow structures with significant density fluctuations disturbing the optical path. Different sources of pressure fluctuations were observed: (1) high-frequency acoustics noise is generated by unsteady flow structures inside the shear layer and (2) an acoustic noise source appears to be located where the unsteady shear layer impinges on the aperture ramp. In addition to these classical acoustic wave-like solutions, two dominant pressure modes were identified by employing proper orthogonal decomposition (POD), one extending over the entire cavity and the other mode displaying raised amplitudes within the observational Nasmyth tube. These pressure POD modes may be a cause for structural vibrations of the telescope.

Figure 2 shows several views of the simulated SOFIA configuration. Note that the aircraft has been simplified slightly, with the most significant simplification being the exclusion of the engines. Simulating the entire flight scale vehicle including geometric complexities and unsteady nonlinear flow physics is a highly complex multi-scale mathematical, physical, and computer science problem. CFD on thousands of cores of a massively parallel supercomputing cluster enable such demanding simulations. Without supercomputers, this analysis at the fidelity level presented herein would not be possible. Strong parallel scaling studies using NASA’s Pleiades supercomputer are presented that show good scaling for the MPI-based numerical method. Dynamic load balancing based on execution time on the individual adaptive mesh refinement (AMR) blocks is used to address irregularities caused by the highly complex geometry. Limits to scaling are identified, and targeted code optimization are discussed.

The current paper is organized as follows: the numerical methods are briefly describe in §2, parallelization strategy in §3, cavity flow physics are analyzed in §4.1-§4.3, and strong scaling performance studies on 400-32000 cores and dynamics load balancing are discussed in §4.4. Finally, we conclude the paper with a summary and outlook on future efforts in §5.

## 2. NUMERICAL METHODS

The LAVA CFD solver [14] was used for this SOFIA application. The LAVA solver is highly flexible with respect to the computational mesh, and supports block-structured Cartesian grids with Adaptive Mesh Refinement (AMR) and immersed boundary capabilities, structured curvilinear, unstructured arbitrary polyhedral, and overset grid coupling. Often, the unsteady compressible Navier-Stokes equations are solved using second-order accurate finite-volume formulations. Higher-order accurate finite-difference discretizations with shock-capturing capabilities are also supported for structured Cartesian grids [5] within LAVA. For this work we utilized the WENO-5Z scheme [5]. Standard and strong stability preserving higher-order Runge-Kutta schemes are available within the code for explicit time integration (the code also supports implicit time integration schemes). For this work we utilized the classical explicit fourth-order Runge-Kutta time integrator.

For Cartesian meshes involving complex geometry, a block-structured, Immersed Boundary (IB) AMR method is uti-

lized [14]. This methodology is capable of automatically generating, refining, and coarsening nested Cartesian volumes given a closed surface triangulation, and hence offers the ability to dynamically track important flow features as they develop. In adaptive methods, one adjusts the computational effort locally to maintain a uniform level of accuracy throughout the problem domain. Cartesian AMR is a proven methodology for multi-scale problems, with an extensive existing mathematical and software knowledge base [1, 2, 3, 4, 24]. The LAVA code leverages data structures and inter-level AMR operators from the high-performance parallel Chombo AMR library [8]. AMR allows the simulation of a wide range of spatial and temporal scales through local refinement. For this work, the constraint that all AMR levels are advanced with the same time step is imposed; i.e., subcycling the levels as is done in the literature [1, 16] is not utilized. A more complete description of the LAVA solver is available in the literature [14, 5, 22].

In the runs, refinement targeted geometric features and sufficient cavity shear layer resolution, with factors of four refinement used between levels. In the code the domain is decomposed into refinement levels of uniform Cartesian mesh, where each level is decomposed into equal sized boxes which are distributed among the MPI ranks.

### 3. PARALLELIZATION STRATEGY

The parallelization strategy currently relies on a pure MPI-based communication pattern based on a SPMD paradigm. The flow domain is decomposed into a hierarchy of refinement levels, and on each level the domain is decomposed into boxes. We impose that the boxes are defined on a box-index-space lattice for each level. Boxes are distributed to processors using a weighted locality preserving knapsack algorithm, where weights are either uniform per box, or dynamically computed using cumulative execution time for each box on the lattice. Dynamic load balancing is used to address irregularities in the performance per box, typically caused by the immersed boundary procedure. All boxes contain three ghost cells per side due to the WENO-5Z scheme. At box interfaces for each level, we do traditional ghost cell exchanges. At coarse-fine interfaces we interpolate to the ghost cells.

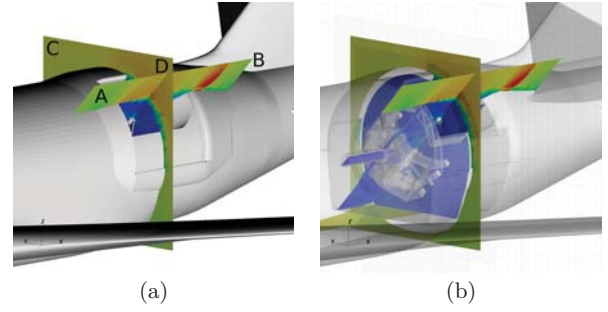
Complex geometry is defined using CAD generated water tight surface triangulations. Triangulations are read from disk at the start of the simulation by each MPI rank, and multi-resolution distance function and binning algorithms are utilized for distance and in/out queries required by the immersed boundary method. Each MPI rank contains the data necessary to independently respond to geometry queries, resulting in a completely parallel immersed boundary method.

## 4. RESULTS

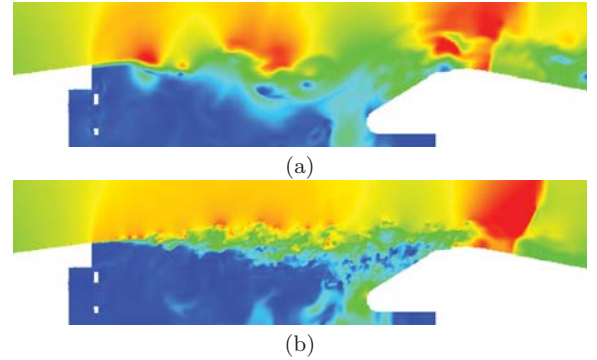
### 4.1 Cavity Flow Physics

CFD based analysis is now presented with the aim of identifying and quantifying acoustic noise generation mechanisms, including shear layer, flow impingement, and modal analysis. The unsteady pressure field in and around the cavity is analyzed by using disturbance flow fields, spectral analysis, and proper orthogonal decomposition (POD).

Figure 3 defines two planar slices used in the remainder of



**Figure 3: Two slices: A-B tilted at 45 degrees, and C-D whose normal is parallel to freestream flow direction. Slices are colored by Mach number used for post-processing analysis. (a) and (b) are of the same view. (a) shows opaque walls, while (b) shows transparent walls.**



**Figure 4: Zoom in on slice A-B as in Figure 3 colored by Mach number showing sensitivity of flow to mesh resolution. (a) coarse grid, (b) fine grid.**

the paper: A-B tilted at 45 degrees, and C-D whose normal is parallel to freestream flow direction. The grid topology is shown in thin gray lines. Most grid points are clustered in the shear layer region where the flow unsteadiness originates. The effect of the grid resolution on the flow field is demonstrated in Figures 4a and 4b.

Mach number contours for the coarse mesh Figures 4a (~170 million cells), and fine mesh Figure 4b (~300 million cells) visualize the transonic flow field around the cavity opening characterized by the highly unsteady shear layer and the interaction of the shear layer with the aft aperture as well as with a re-compression shock located just downstream of the cavity opening. The simulation with refined mesh (Figure 4b) shows an earlier breakup of the shear layer and finer scale vortical structures in comparison to the coarse grid simulation (Figure 4a). The interaction between the unsteady shear layer and the re-compression shock is reduced as a consequence of the smaller unsteady vortical flow structures.

The time-averaged colored streamwise velocity contours and solid pressure contour lines for the fine and coarse grids can be compared in Figure 5. At first glance, the general mean



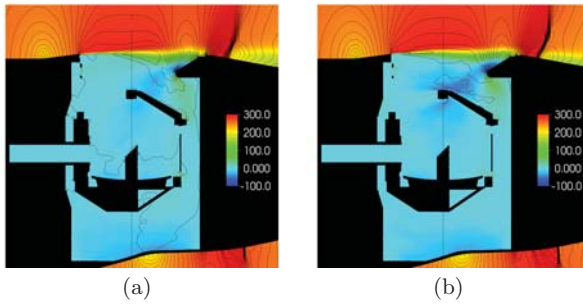


Figure 5: Mean flow slices A-B as in Figure 3 colored by streamwise velocity, with pressure contour lines showing sensitivity to mesh resolution. (a) coarse grid, (b) fine grid.

flow field features appear very similar for both simulations. A closer look reveals some differences in the mean flow distortions of the velocity and pressure fields. The mean pressure distribution along the shear layer is slightly different, and the downwash at the shear layer impingement point at the aft aperture is increased for the fine grid simulation. Moreover, the fine grid allows for a better resolution of the shock structure.

The iso-contours of vorticity in Figure 6 provide an overview of the general features and complexity of the flow field. At the inception of the shear layer, the spanwise roll-up of the vortex sheet, characteristic of the primary instability, i.e., Kelvin-Helmholtz instability, can be observed. Shortly downstream a spanwise modulation occurs introducing 3D flow structures. Subsequent breakdown of coherent structures results in fine scale turbulent-like flow features. Note that at these flow conditions a grid resolution for direct numerical simulations (DNS), where all temporal and spatial scales are resolved down to the Kolmogorov length scale, would require orders of magnitude more grid points.

## 4.2 Spectral Analysis

In contrast to small differences in mean flow distortion, the differences in the unsteady part of the solution  $u' = u - \bar{u}$  is very significant. Figure 7 shows the breakdown of large coherent structures on the fine mesh reduces the total power spectral density (PSD) values. It is well known that not resolving the unsteady flow features will lead to highly energetic spanwise coherent flow structures. This strongly coherent flow field tends to cause elevated acoustic noise levels in and around the cavity. The fine grid solution provides an improved mechanism for the energy cascade where the spanwise coherence is reduced and fine scale turbulent like flow features can be observed. It must be pointed out that telescope image distortion depends strongly on the actual shear layer thickness. In the large eddy simulation, the spreading rate of the shear layer in turn depends strongly on the eddy viscosity. Unsteadiness in the flow field needs to be well resolved to accurately model the shear layer mixing and provide precise predictions of the eddy viscosity. As a side note, Reynolds-averaged Navier-Stokes (RANS) simulation strategies are computationally far less expensive than the current ILES approach. RANS approaches perform very well in predicting eddy viscosities for attached wall bounded

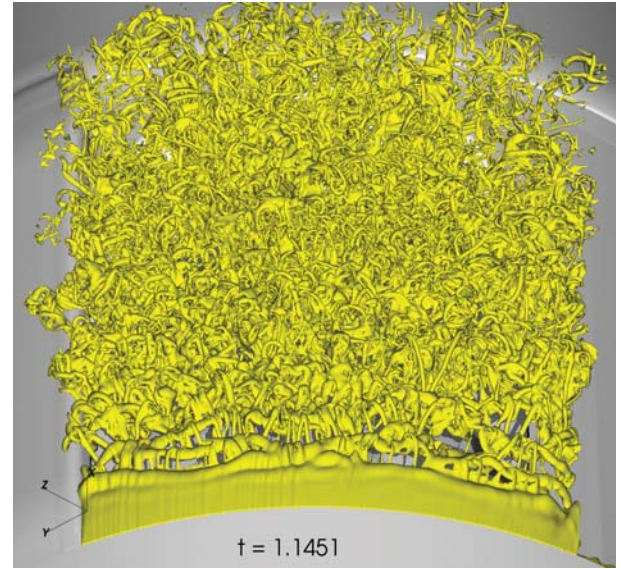


Figure 6: Vorticity contours in the shear layer at the cavity opening.

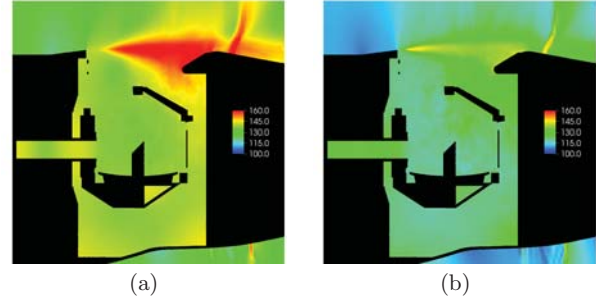


Figure 7: Slices Slice A-B as in Figure 3 colored by total power spectral density (PSD) [dB] showing sensitivity of mean flow to mesh resolution. (a) coarse grid, (b) fine grid.

flows but for separated flows and shear layers, reliable data cannot be expected. This is why we chose the more expensive ILES approach. In addition, RANS approaches do not provide insight about the occurrence of unsteady cavity modes discussed below.

To assess the grid resolution in this simulation, the power spectrum of the streamwise disturbance velocity  $u'/u_{ref}$  was computed at discrete points in the shear layer, i.e., points 0-3 in Figure 8a. The high frequency peak in the spectrum at around 1000 Hz at the inception of the shear layer (point 0) is associated with primary Kelvin-Helmholtz instability. The spectrum quickly broadens and increases in amplitude due to the rapid transition process. The blue dashed line marks the theoretical sub-inertial scaling law that appears to be observed for the high-frequency content of the disturbance velocity signal. The unsteady shear layer causes high pressure fluctuation away from the the shear layer. The Fourier spectrum of pressure for points 0-2 and points 6-8

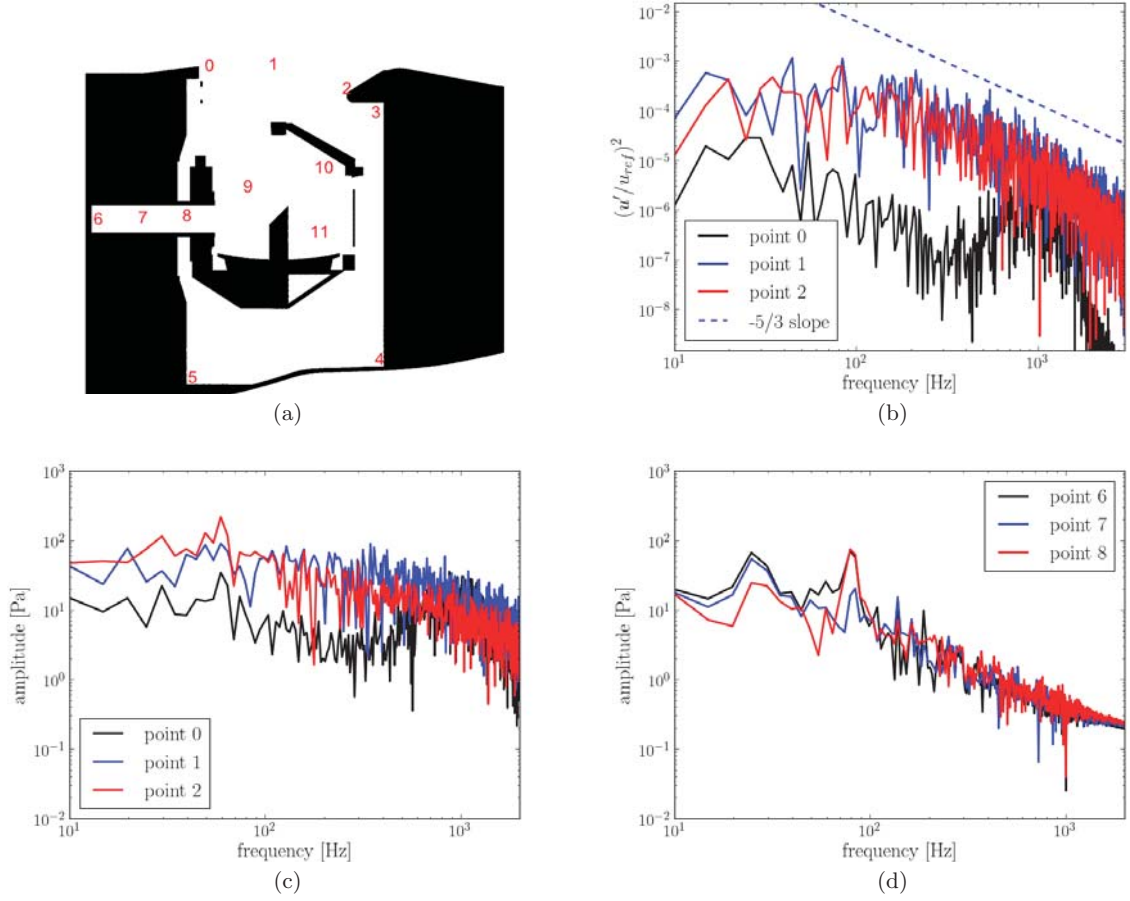


Figure 8: (a) Sample point locations shown on slice A-B. Shear layer point spectra: (b) streamwise velocity power spectra at points 0-2, and (c) pressure spectra amplitude for points 0-2 and (d) points 6-8.

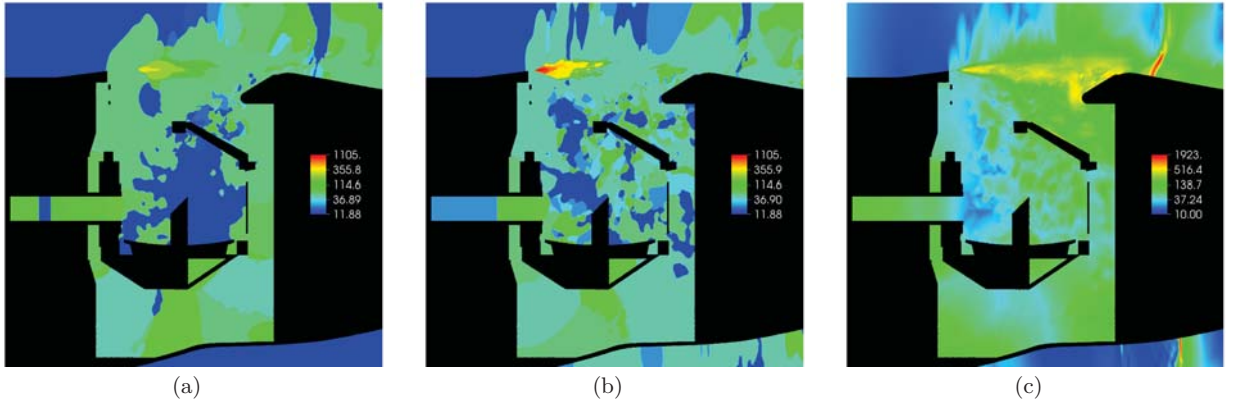
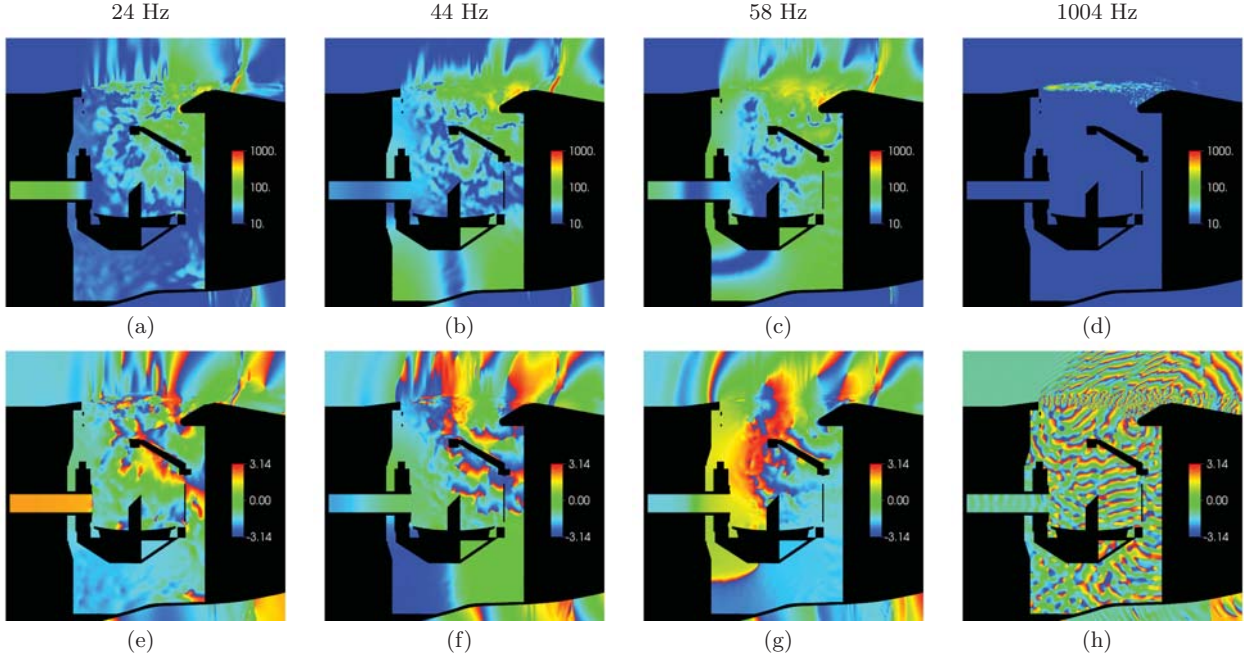


Figure 9: Slices A-B as in Figure 3 colored by spectral frequency [Hz] at peak amplitude showing sensitivity of mean flow to mesh resolution. (a) coarse grid, (b) fine grid. (c) Fine grid slice colored by peak spectral amplitude, showing locations with dominant frequencies.



**Figure 10: Phase analysis for pressure at different frequencies. Top row: spectral amplitude [Pa] at (a) 24 Hz (b) 44 Hz (c) 58 Hz (d) 1004 Hz; Bottom row: phase plots [rad] at (e) 24 Hz (f) 44 Hz (g) 58 Hz (h) 1004 Hz.**

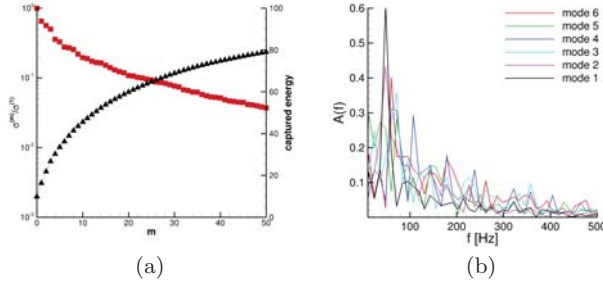
are shown in Figures 8c and 8d, respectively. In addition to the information from the power spectra in Figure 8a, the pressure Fourier spectra for point 1 and point 2 show peaks at around 57 Hz. The additional pressure sensors at probe locations 6-8 are located within the Nasmyth tube. Two peaks at approximately 20Hz and 90Hz are identified. Further analysis elaborates on particular frequency characteristics, i.e., broad vs. narrow band. In contrast to the velocity power spectra inside the shear layer which are broad band, the shear layer impinging on the aft aperture and some geometric features of the cavity appear to introduce narrow band frequency characteristics.

The spectral characteristics of the unsteady flow and the origin of these peaks in the pressure Fourier spectra is analyzed by applying discrete Fourier transforms in time to the data collected on the two sampling planes shown in Figure 3. Figures 9a and 9b display the frequency at the peak amplitude in the Fourier spectrum for the coarse and fine grids. The highest frequency at the inception of the shear layer is associated with the primary instability. Even though the boundary layer at the beginning of the cavity opening is turbulent, a “secondary” transition process occurs when the shear layer forms [19, 20]. The coarse grid simulation fails at predicting the initial growth rate of the shear layer instability and its subsequent breakdown to turbulence. For the coarse grid simulation, the peak frequency is significantly lower and the peak occurs further downstream. Thus, the shear layer breakdown process is delayed. Nearly uniform low peak frequencies are observed everywhere inside the cavity. The acoustic near field of the shear layer shows larger peak frequencies for the coarse grid simulation in comparison to the fine grid results. The peak amplitudes in Figure

9c are largest at the shocks and in the shear layer region. The peak amplitude inside the cavity is an order of magnitude lower than in the shear layer but significant enough to cause small vibrations of telescope holding structure. Moreover, these pressure fluctuation may be close to a resonance frequency of the telescope structure.

Figure 10 presents temporal Fourier analysis results. Amplitude and phase plots in plane A-B are shown for selected frequencies. At 24Hz large pressure amplitudes can be observed in the Nasmyth tube. The phase plot indicates that the pressure oscillation is in phase throughout the tube. The standing wave with a frequency of 24Hz is perfectly tuned to the length of the tube. For all three lower frequencies (24Hz-58Hz), no wave-like phase pattern can be observed due to the large wave lengths at these frequencies. At 44Hz and 58Hz, large amplitudes can be observed inside the cavity. The phase plots are not straight forward to interpret since the cut planes only provide a 2D view of a 3D spectral mode. For 44Hz, there appears to be a phase shift from the lower left corner to the lower right corner. This pressure mode can induce a left-to-right rocking motion of the telescope holding structure that must be controlled with structural damping. The interaction of the shear layer with the aft aperture is picked up in the 44Hz and 58Hz modes. The phase plot for 44Hz also points towards a source location where the shear layer impinges onto the end of the cavity opening. The iso-contour lines of constant phase display a radial pattern away from the source location. At high frequency 1004Hz, the highest fluctuation amplitude is located in the shear layer. The frequency of 1004Hz was chosen since the most amplified initial disturbance growth was obtained around this particular frequency. Initially a very coherent





**Figure 11: (a) Singular value of  $m^{th}$  POD mode normalized with maximum singular value  $\sigma_1$  (red triangles) and total energy captured by the first  $N$  POD modes (black squares). (b) Discrete Fourier transform of the POD time-coefficients for the six most energetic POD modes. Results shown for plane A-B as defined in Figure 3.**

pattern can be identified inside the shear layer by the planar wave fronts in the phase plot. The short wave-length is related to the convection speed and high frequency. Various acoustic wave patterns are visible in the phase plots. The large unsteady coherent flow structures generate acoustic noise which propagates away from the shear layer. The sonic line separates the silent region upstream of the beginning of the cavity from the acoustic near field of the shear layer and the cavity flow. The interaction of the shear layer with aft aperture also generates some high frequency pressure oscillations. The acoustic wave-like phase pattern quickly loses amplitude as the waves are reflected and interfere with another inside the cavity.

### 4.3 Proper Orthogonal Decomposition

POD results in a decomposition of the flow field into a set of basis functions that capture most of the flow energy as defined by a user-defined norm with the least number of modes [15]. In many fluid dynamics applications, POD is used to identify energetic flow structures. It must be noted that it is not guaranteed that existing POD modes have dynamic significance. Freund and Colonius [12] applied POD to a Mach 0.9 turbulent jet. They note that, for example, the radiated sound by turbulence amounts to a very small fraction of the total energy. Thus, a POD approach in the temporal domain may not be able to capture the physics of the radiated sound waves. Furthermore, it is not necessarily clear what the appropriate norm for capturing particular flow physics is. Freund and Colonius [12] have found that the efficiency at representing the initial flow data depends strongly upon the norm used and what data are represented. For problems where the spatial degrees of freedom is much greater than the number of available time-steps,  $N_{tot}$ , it is more efficient to employ the “snapshot” method [21] for POD. For more details about the POD snapshot method used here, see Chatterjee [7]. For a basic introduction to POD, see [12, 18] on details on how it can be applied to fluid dynamics problems.

The flow field can be reconstructed/recovered from the eigen-

modes and time coefficients of the POD modes:

$$\vec{q}(\vec{x}, t) \approx \sum_{n=0}^I a^{(n)}(t) \vec{\psi}^{(n)}(\vec{x}). \quad (1)$$

The POD modes are orthogonal to each other and sorted by their respective energy norm contents, whereby the drop-off in mode energy toward the higher mode numbers is typically significant. Therefore, a small number of modes,  $I$ , will suffice for capturing a large percentage of the flow’s energy. This is particularly the case when the flow dynamics are governed by large energetic structures of organized (or coherent) fluid motion (such as in pure 2D simulations).

In the current study, the snapshot method was employed and the energy norm is based on the static pressure. In order to measure the coherence in the flow field we introduce the factor  $\psi_N$ . This factor provides a measure for how much energy in the pressure field is captured by the first  $N$  POD modes:

$$\psi_N = \left( \sum_{m=1}^N \sigma_m \right) / \left( \sum_{m=1}^{N_{tot}} \sigma_m \right), \quad (2)$$

where  $\sigma_m$  are the singular values corresponding to the  $m^{th}$  POD mode,  $N_{tot}$  is the total number of POD modes (equal to the number of time-steps) and  $N < N_{tot}$ . The singular value is related to the energy of the POD defined by the user specified norm. A significant drop off in the singular values is expected for highly coherent flows. The POD analysis was applied to the entire A-B planar slice, and therefore, a large amount of captured energy is associated with the highly unsteady shear layer region. The energy distribution for the POD modes is shown in Figure 11a. As demonstrated for the power spectra of streamwise velocity in Figure 8b, the energy in the flow is distributed over a broad turbulent like spectrum. Based on the energy distribution of the discrete POD modes, it appears that the flow field is not very coherent. To capture 60-80 per cent of the energy 30-50 POD modes are required. For the coarse grid calculation, the number of POD modes required to capture the same amount of energy is significantly less because the coarse grid is not able to support the short wave lengths. Thus, the coarse grid calculation tends to over-predict the coherence in the flow field. The Fourier amplitudes of the time coefficients indicate that the temporal evolution for some of the POD modes is dominated by discrete frequencies in the spectrum. A large peak in Fourier spectra at around 60Hz is apparent for the time coefficients of POD modes 1 and 2. The POD modes 3 through 6 show a large peak around 70-80Hz. For POD mode 4, a sequence of discrete peaks is noted in the Fourier spectrum that is associated to an interaction between different temporal modes. Higher and lower harmonics can be generated by a mode-mode interaction in the frequency domain.

The spatial distributions of the POD modes are illustrated in Figures 12a-d. It must be noted that the scaled POD mode shapes in Figures 12a-d do not provide information about amplitude of the pressure oscillations one may expect in particular parts of the flow field. As expected, the POD mode shapes clearly highlight the shear layer as the region where most of the energy is contained. POD mode 1 provides strong evidence of a coupling between the shear layer

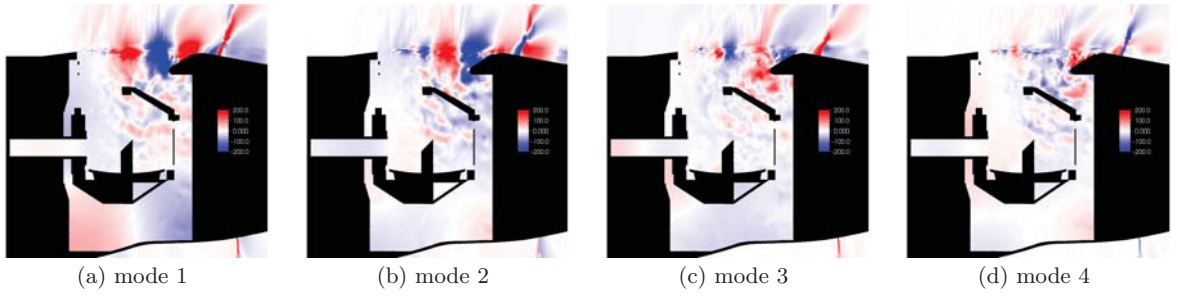


Figure 12: (a)-(d) Four most energetic POD modes 1-4.

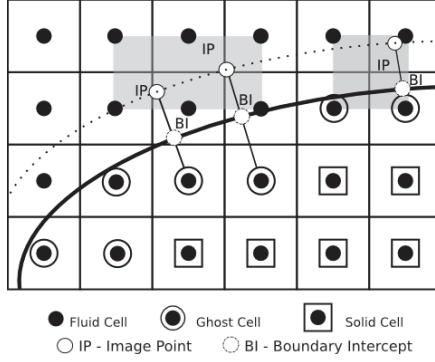


Figure 13: Immersed ghost-cell boundary treatment as described in Kiris *et al.* [14]

interaction with the aft aperture and the cavity mode. The modulation in POD mode 1 in the shear layer has a dominant wave length of approximately half of the opening of the cavity and the wave length inside the cavity is equal to the length of the cavity. The interaction of the unsteady shear layer with the re-compression shock is also picked up by the POD modes. No upstream effect of this interaction is expected due to the supersonic flow upstream of the re-compression shock. POD modes 2-4 appear to mainly highlight the interaction between the shear layer and the aft aperture. Some less significant coupling between the activity in the shear layer and pressure fluctuations inside the cavity can also be noted for POD modes 2-4. The downwash at the aft aperture that provides momentum flux inside the cavity is present in all four POD modes. The downwash can strongly affect the clarity of the telescope images because it crosses the optical path and more importantly it induces structural vibrations of the telescope arm.

#### 4.4 Solver Parallel Performance

The analysis of the flow field in the previous section has highlighted the importance of providing sufficient grid resolution such that the wide range of physically relevant temporal and spatial scales can be resolved. Even with very efficient higher-order numerical schemes, significant resources are required to conduct such high fidelity CFD analysis. In the current section, we want to demonstrate two key capabilities of the LAVA solver: (1) scalability can be achieved for up to at least tens of thousands of processors and (2) the ability

to dynamically adjust to the non-uniform per box timings of the immersed boundary method. To address the first objective, strong scaling studies were performed for both fine and super-fine meshes with the highly complex SOFIA geometry with multi-levels of resolution. All runs were conducted using NASA's Pleiades supercomputer. The Pleiades nodes used are composed of Intel Xeon E5-2680v2 processors (Ivy Bridge). The Ivy Bridge memory hierarchy is as follows: L1 instruction cache: 32 KB, private to each core; L1 data cache: 32 KB, private to each core; L2 cache: 256 KB, private to each core; L3 cache: 25 MB, shared by 10 cores in each socket; and DDR3 FB-DIMMs main memory: 32 GB per socket, total of 64 GB per node. Sockets are connected with a 32 GB/s link. Each node (20 cores) was assigned to 20 MPI ranks, and internode communication is performed over InfiniBand, with all nodes connected in a partial 11D hypercube topology.

The fine mesh strong scaling study was performed using three different box sizes,  $8^3$ ,  $16^3$ , and  $32^3$  and 7 levels of refinement, and 400-8000 cores. At  $8^3$  the tagging resulted in 392,461 total boxes or 200,940,032 cells, and 1,076,912,984 total cells including 3-ghost cells on all sides of the box (i.e.  $14^3$ ). At  $16^3$  the tagging resulted in 59,904 total boxes or 245,366,784 cells, and 637,857,792 total cells including 3-ghost cells on all sides of the box (i.e.  $22^3$ ). At  $32^3$  the tagging resulted in 10,362 total boxes or 339,542,016 cells, and 568,583,664 total cells including 3-ghost cells on all sides of the box (i.e.  $38^3$ ). At  $32^3$  there are not enough boxes / work for the large processor count.

The top row of Figure 14 presents the results of the strong scaling study on the fine mesh. Three components of the solver were characterized, (a) entire Runge-Kutta time integration step including communication, (b) WENO-5Z flux and divergence operators, and (c) immersed boundary treatment. The entire Runge-Kutta time-integration step includes four substeps for the fourth-order accurate time integrator. At each substep, the solver proceeds as follows: first, the convective fluxes are first computed at each node utilizing Lax-Friedrichs flux vector splitting. The fluxes at the half points are then obtained via WENO reconstruction from the fluxes at nodes. The second-order viscous terms are evaluated in conservative form where we obtain the viscous fluxes at the half points employing classical finite difference evaluation. Finally, at each computational cell the discrete divergence operator is applied to the convective and viscous fluxes previously obtained at the half-points. After each



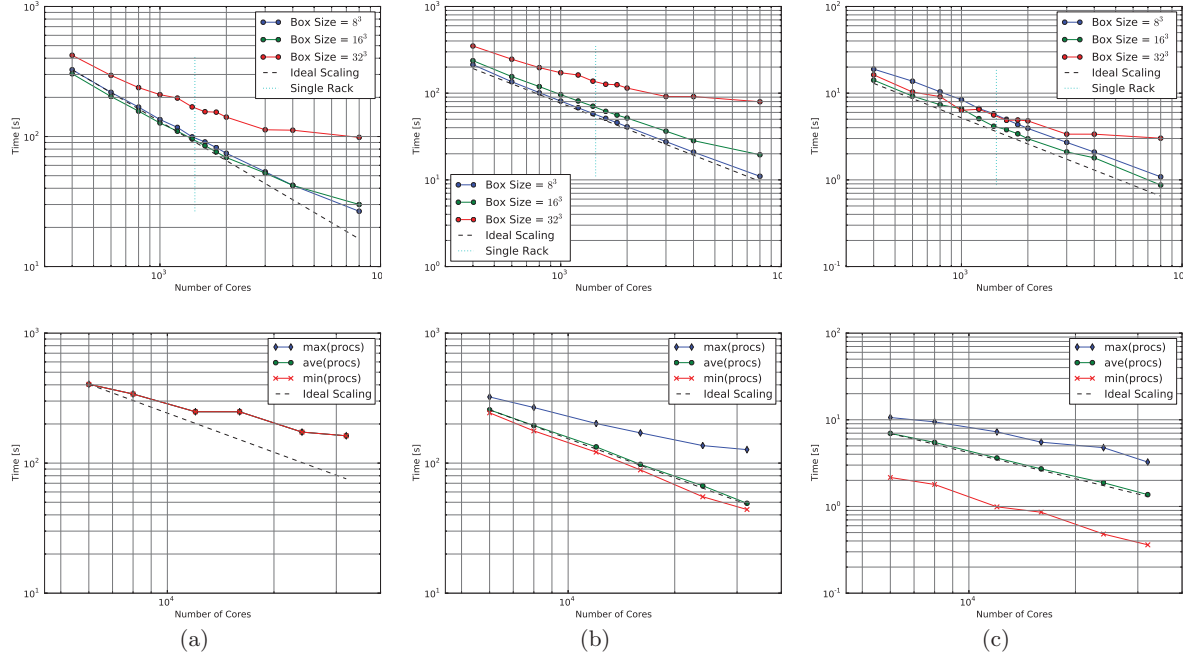


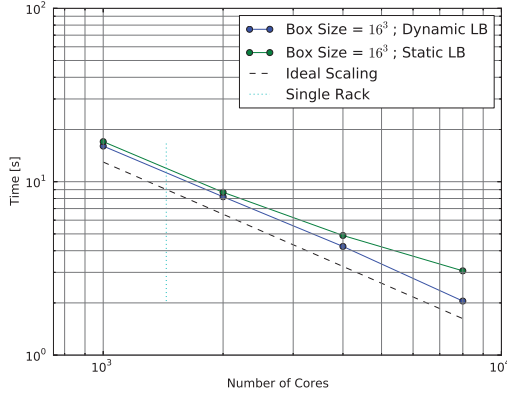
Figure 14: Strong parallel scaling study for 25 time steps: top row is  $\sim 300$  million cells (fine mesh), max time over all procs shown; bottom row is 4.4 billion cells (super-fine mesh), min/ave/max time over all procs shown for box size of  $32^3$ . (a) Entire RK-solve including communication, (b) WENO-5Z flux and divergence operators, (c) immersed boundary conditions.

step, block, coarse-fine interface, and domain ghost cells need to be filled. For coarse-fine grid interfaces an interpolation procedure is needed to fill the coarse and fine grid neighboring ghost cells. The WENO-5Z flux reconstruction procedure consists of three basic steps: (1) forward characteristic transform of convective flux, (2) computation of smoothness indicator and nonlinear weights for three candidate stencils, and (3) backward characteristic transform. The WENO-5Z flux reconstruction is applied to individual grid lines to keep the memory local in cache. The immersed boundary method is extremely powerful because it provides the capability of computing flows around highly complex geometries without significant manual mesh generation. From a programming perspective, it is, however, a non-uniform part of the code. The current implementation of the ghost cell method is first-order accurate at the wall. Each ghost-cell in Figure 13 needs to be visited and an immersed boundary ghost cell procedure is applied. For higher-order accurate immersed-boundary schemes, as the one currently being developed in LAVA [6], the cost of the immersed boundary treatment increases. Hence, processors with immersed boundary ghost cells or irregular points incurs a larger overhead. In order to load balance the work on the processors, we utilize dynamic load balancing where the block assignment per processors is re-evaluated after a number of time-steps based on recorded execution times for each box, where the boxes are on a lattice based uniform box-index-space per level.

The fine mesh scaling study shows that  $8^3$  boxes are preferred for the WENO-5Z flux and divergence operators (Fig-

ure 14b top row) due to locality in the memory hierarchy. This trend is not the same for the immersed boundary conditions (Figure 14c top row), where  $16^3$  is preferred. Note the  $32^3$  flat-line in Figure 14c top row, is due to not enough boxes/work available. Figure 14a top row, which shows the performance for all operations during 25 Runge-Kutta steps scales very well for the smaller sized boxes. No significant dependence was found when more than one rack was used (dotted blue line in Figures).

A super-fine mesh was generated by refining the shear layer region with an additional refinement level, and a strong scaling study was conducted using 6000-32,000 cores. For this super-fine mesh, smaller sized boxes were attempted ( $8^3$  and  $16^3$ ), which would yield more boxes per level and improved load balancing, but memory was not sufficient due to the volume mesh meta-data [23]. Mesh meta-data compression and/or hybrid MPI/OpenMP (anticipated future work) should allow for smaller boxes, and therefore improved scaling. Boxes were fixed at  $32^3$  and the tagging resulted in 136,603 total boxes or 4,476,207,104 cells, and 7,495,679,816 total cells including 3-ghost cells on all sides of the box (i.e.  $38^3$ ). Figure 14 bottom row presents the strong scaling study for the 8 level mesh. Similar trends are observed as compared to the fine mesh study. Due to the multi-scale nature of the refinement hierarchy there are many more boxes on the finer levels for this application. The 8 level computations have the following number of boxes for each level: [8, 8, 12, 52, 474, 6623, 3660, 125766]. In the current algorithm, these boxes are distributed to processors per



**Figure 15: Strong parallel scaling study for 100 time steps, adaptively regridding every 5 steps, comparing dynamic and static load balancing schemes. Simulation is for fine mesh, with max time over all procs shown. Timings are for immersed boundary conditions routines only, using box size of  $16^3$ .**

level, which for portions of the algorithm can lead to inter-level load imbalances, as seen in Figure 14 bottom row, when the number of boxes per level is non-uniform, e.g. 8 boxes on coarsest level vs 125766 on finest level.

On the fine mesh, another strong scaling study was performed illustrating the impact of timer based “dynamic” load balancing during adaptive mesh refinement as compared to uniform or “static” loads. Dynamic load balancing is based on timers for each box in the AMR hierarchy, weighting box loads by the accumulated time. This is possible due to the fixed-box-size box layout used in these simulations, where the boxes are stored on a lattice in a box index-space for each level. Dynamic load balance timings were recorded for all box operations. Static load balancing is based on a uniform weighting of the box loads. For this scaling study the simulation was run for 100 time steps, with re-gridding and load balancing occurring every 5 steps. Figure 15 shows an improvement for the immersed boundary condition routine due to dynamic loads. No significant improvement was observed for the remainder of the routines as the load is well balanced with naive uniform load balancing (see Figure 14), given a sufficient workload.

## 5. CONCLUSION

The paper presents a higher-order implicit large eddy simulation of a large telescope mounted in a large open cavity in the aft of a Boeing 747SP. Employing conventional CFD approaches, significant time must be spent to generate a computational mesh which allows for high quality CFD results. The current approach which utilizes an automatically generated block-Cartesian data layout, allows for fast data access, and which promises fast execution times.

Different grid resolutions were used to study sensitivities of the solution with respect to the computational grid. On the fine grid, theoretical inertial sub-range scaling was obtained

in the power spectra of streamwise and spanwise velocities within the shear layer region. The unsteady pressure field in and around the cavity was analyzed by using spectral analysis and proper orthogonal decomposition (POD).

Two important noise generation mechanisms were identified: (1) high-frequency acoustics noise is generated by unsteady flow structures inside the shear layer and (2) an acoustic noise source appears to be located where the unsteady shear layer impinges on the aperture ramp. In addition to these classical acoustic wave-like solutions, two dominant pressure modes were identified, one extending over the entire cavity and the other pressure mode displaying raised amplitudes within the Nasmyth tube. The cavity mode frequency is approximately 60Hz. A sound pressure level of 120-130dB was determined at the rear wall of the cavity. These values are consistent to the results reported by Cummings [9]. A strong interaction between the cavity mode and the unsteadiness of the shear layer was shown in the POD analysis. The wavelength of the shear layer POD mode which is coupled to the cavity mode is significantly larger than that of the primary shear layer instability. In contrast to the primary shear layer instability wavelength, which scales with the momentum thickness of the shear layer, this wavelength or its higher harmonics appear to scale with the distance of the shear layer inception point to the point where the shear layer impinges on the cavity opening. Hence, this distance will determine possible wavelengths in the shear layer which couple to the cavity mode. By employing this underlying physical mechanism the characteristics of the cavity can be influenced by means of active or passive flow control. Furthermore, the shear layer thickness can be effected by upstream perturbations. In addition to the cavity mode, a downwash [10], i.e., momentum flux into the cavity, was observed to originate at the location where the shear layer impinges on the cavity opening. The downwash is aimed in the direction of the telescope and may induce unsteady loading and perturbations to the optical path.

In the solver performance study we were able to demonstrate that LAVA is scalable to at least 32,000 processors for this highly complex geometry, multi-scale problem. Moreover, we also demonstrated the ability of LAVA to dynamically adjust to the non-uniform loads induced by the immersed boundary method. The scaling study revealed several areas for improvement. Our parallel optimization plan for the future is to adopt a hybrid MPI/OpenMP approach where boxes are allocated to an MPI rank and dynamically distributed to OpenMP threads. This approach should improve load imbalances due to the immersed boundary condition in addition to memory limits depending on MPI rank (e.g. volume and surface mesh related). Another limit to scaling, identified by the scaling study, that we will pursue, is to reduce the memory overhead of storing the volume mesh meta-data on each processor as in [23]. Further performance optimizations will include a focus on increasing percent of nominal peak performance [13, 17], within the immersed boundary and block-structured AMR methodologies [14].

## 6. ACKNOWLEDGMENTS

We gratefully acknowledge the SOFIA project for access to CAD models and conditions. C. Brehm supported under contract STC-NNA10DF26C with NASA.

## 7. REFERENCES

- [1] A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, and M. L. Welcome. A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *J. Comp. Phys.*, 142:1–46, 1998.
- [2] M. F. Barad and P. Colella. A Fourth-Order Accurate Local Refinement Method for Poisson’s Equation. *J. Comp. Phys.*, 209(1):1–18, October 2005.
- [3] M. F. Barad, P. Colella, and S. G. Schladow. An Adaptive Cut-Cell Method for Environmental Fluid Mechanics. *Int. J. Numer. Meth. Fluids*, 60(5):473–514, 2009.
- [4] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.*, 82(1):64–84, May 1989.
- [5] C. Brehm, M. Barad, J. Housman, and C. Kiris. A Comparison of Higher-Order Shock Capturing Schemes Within the LAVA CFD Solver. In *AIAA Paper 2014-1278, Jan 13-17, National Harbor, Maryland*, 2014.
- [6] C. Brehm and H. Fasel. A novel concept for the design of immersed interface methods. *Journal of Computational Physics*, 242(0):234 – 267, 2013.
- [7] A. Chatterjee. An introduction to the proper orthogonal decomposition. *Current science*, 78(7):808–817, 2000.
- [8] P. Colella, D. T. Graves, T. J. Ligocki, D. F. Martin, D. Modiano, D. B. Serafini, and B. V. Straalen. Chombo Software Package for AMR Applications - Design Document. unpublished, 2000.
- [9] S. B. Cumming, M. Smith, L. Cliatt, M. Frederick, and F. by Institutions. Aerodynamic and acoustic flight test results for the stratospheric observatory for infrared astronomy. 2013.
- [10] C. Engfer, T. Lutz, and E. Krämer. Characterization of the cavity shear layer of the stratospheric observatory for infrared astronomy by means of pressure sensor data and a hybrid rans-les study. 2013.
- [11] C. Engfer, E. Pfüller, M. Wiedemann, J. Wolf, T. Lutz, E. Krämer, and H.-P. Röser. Evaluation of the aero-optical properties of the sofia cavity by means of computational fluid dynamics and a super fast diagnostic camera. In *SPIE Astronomical Telescopes + Instrumentation*, pages 844412–844412. International Society for Optics and Photonics, 2012.
- [12] J. B. Freund and T. Colonius. Pod analysis of sound generation by a turbulent jet. In *AIAA 2002-0072, 40th AIAA Aerospace Sciences Meeting and Exhibit*, pages 1–9, 2002.
- [13] B. Hejziahosseini, D. Rossinelli, C. Conti, and P. Koumoutsakos. High throughput software for direct numerical simulations of compressible two-phase flows. In *High Performance Computing, Networking, Storage and Analysis (SC), 2012 International Conference for*, pages 1–12. IEEE, 2012.
- [14] C. Kiris, M. Barad, J. Housman, E. Sozer, C. Brehm, and S. Moini-Yekta. The LAVA Computational Fluid Dynamics Solver. In *AIAA Paper 2014-0070, Jan 13-17, National Harbor, Maryland*, 2014.
- [15] J. Lumley. The structure of inhomogeneous turbulent flows. *Atm. Turb. And Radio Wave Prop.*, Nauka, Moscow and Toulouse, France, eds. Yaglom and Tatarsky, pages 166–178, 1967.
- [16] D. F. Martin, P. Colella, and D. Graves. A Cell-Centered Adaptive Projection Method for the Incompressible Navier-Stokes Equations in Three Dimensions. *J. Comput. Phys.*, 227:1863–1886, 2008.
- [17] D. Rossinelli, B. Hejziahosseini, P. Hadjidoukas, C. Bekas, A. Curioni, A. Bertsch, S. Futral, S. J. Schmidt, N. A. Adams, and P. Koumoutsakos. 11 pflop/s simulations of cloud cavitation collapse. In *Proceedings of SC13: International Conference for High Performance Computing, Networking, Storage and Analysis, SC ’13*, pages 3:1–3:13, New York, NY, USA, 2013. ACM.
- [18] C. W. Rowley. *Modeling, simulation, and control of cavity flow oscillations*. PhD thesis, California Institute of Technology, July 2001.
- [19] M. Shur, P. Spalart, and M. Strelets. Noise Prediction for Increasingly Complex Jets. Part 2: Applications. *International Journal of Aeroacoustics*, 4(4):247–266, 2005.
- [20] M. Shur, P. Spalart, and M. Strelets. Noise Prediction for Increasingly Complex Jets. Part 1: Methods and Tests. *International Journal of Aeroacoustics*, 4(3):213–246, 2005.
- [21] L. Sirovich. Turbulence and the Dynamics of Coherent Structures. *Quarterly of Applied Mathematics*, XLV(3):561–590, 1987.
- [22] E. Sozer, C. Brehm, and C. Kiris. Gradient Calculation Methods on Arbitrary Polyhedral Unstructured Meshes for Cell-Centered CFD Solvers. In *AIAA Paper 2014-1440, Jan 13-17, National Harbor, Maryland*, 2014.
- [23] B. Van Straalen, P. Colella, D. T. Graves, and N. Keen. Petascale block-structured amr applications without distributed meta-data. In *Euro-Par 2011 Parallel Processing*, pages 377–386. Springer, 2011.
- [24] Q. Zhang, H. Johansen, and P. Colella. A fourth-order accurate finite-volume method with structured adaptive mesh refinement for solving the advection-diffusion equation. *SIAM Journal on Scientific Computing*, 34(2):179–201, 2012.